# Finding faint H I structure in and around galaxies: scraping the barrel

D. Punzo[a,*], J.M. van der Hulst[a], J.B.T.M. Roerdink[b]

[a] *Kapteyn Astronomical Institute, University of Groningen, Landleven 12, 9747 AD Groningen, The Netherlands*
[b] *Johann Bernoulli Institute for Mathematics and Computer Science, University of Groningen, Nijenborgh 9, 9747 AG Groningen, The Netherlands*

arXiv:1609.03782v1 [astro-ph.IM] 13 Sep 2016

**Abstract**

Soon to be operational H I survey instruments such as APERTIF and ASKAP will produce large datasets. These surveys will provide information about the H I in and around hundreds of galaxies with a typical signal-to-noise ratio of $\sim 10$ in the inner regions and $\sim 1$ in the outer regions. In addition, such surveys will make it possible to probe faint H I structures, typically located in the vicinity of galaxies, such as extra-planar-gas, tails and filaments. These structures are crucial for understanding galaxy evolution, particularly when they are studied in relation to the local environment. Our aim is to find optimized kernels for the discovery of faint and morphologically complex H I structures. Therefore, using H I data from a variety of galaxies, we explore state-of-the-art filtering algorithms. We show that the intensity-driven gradient filter, due to its adaptive characteristics, is the optimal choice. In fact, this filter requires only minimal tuning of the input parameters to enhance the signal-to-noise ratio of faint components. In addition, it does not degrade the resolution of the high signal-to-noise component of a source. The filtering process must be fast and be embedded in an interactive visualization tool in order to support fast inspection of a large number of sources. To achieve such interactive exploration, we implemented a multi-core CPU (OpenMP) and a GPU (OpenGL) version of this filter in a 3D visualization environment (`SlicerAstro`).

*Keywords:* radio lines: galaxies, techniques: image processing, scientific visualization

## 1. Introduction

Radio data are intrinsically noisy and most sources are faint and often extended (see for example the WHISP catalog, van der Hulst et al. (2001)). Very faint coherent H I signals, below a 3 sigma *rms* noise level, are difficult to find (Popping et al., 2012). Depending on the source structure, spatial and/or spectral smoothing can increase the signal-to-noise ratio. Smoothing is usually applied to multiple spatial and spectral scales to ensure that sources of different size are extracted at their maximum integrated signal-to-noise ratio.

In upcoming blind H I surveys such as WALLABY, using the ASKAP telescope (Johnston et al., 2008; Duffy et al., 2012), and the shallow and medium-deep APERTIF surveys, using the WSRT telescope (Verheijen et al., 2009), source finding will be a major concern. Source finders (e.g., Whiting, 2012; Serra et al., 2015) are designed to automatically detect all the sources in the field and to achieve this goal they must employ an efficient mechanism to discriminate between interesting candidate sources and noise. Due to the complex 3-D nature of the sources (Sancisi et al., 2008) and the noisy character of the data, constructing a fully automated and reliable pipeline is not trivial. Popping et al. (2012) reviewed the current state of the art and described the issues connected with the noisy nature of the data, and the various methods and their efficiency.

In the source-finding process, *masks* are generated enclosing the sources. The determination of the final masks involves a variety of filtering operations in order to pick up faint and extended emissions. However, users are ultimately provided with the mask and data products determined from the original data within the masks. In order to examine the original data within and around the mask, to check the performance of the source finding process and to investigate whether all faint structures have been included, it is necessary to have a visualization tool that not only shows the original data and the mask, but also has the ability to interactively filter the data to bring out the very faint structures in the data.

Our goal therefore is the development of a suitable filtering method in a 3D visualization environment that maximizes the local signal-to-noise ratio of the very faint structures (signal-to-noise ratio $\sim 1$) while preserving its specific 3-D structure (e.g. tidal tails, filaments and extra-planar gas). Ideally, the method should be adaptive (in such a way that the user does not have to explore a large parameter space to get the best result), interactive, and fast, i.e. applicable in real-time. In this paper, we explore a number of existing filtering methods in combination with a 3D visualization tool (Punzo et al., 2015) in order to find a method fulfilling such requirements.

*Corresponding author
*Email address:* `D.Punzo@astro.rug.nl` (D. Punzo)

In Section 2 we describe the datasets used for our investigation and in Section 3 we give an overview of state-of-the-art filtering packages and algorithms, with a focus on radio astronomy. We also describe the filtering techniques chosen for the analysis performed in this paper. In Section 4 we report an analysis of the best parameters for each of the filtering methods. In Sections 5 and 6 we test the quality and the performance of the filtering algorithms implemented. In Section 7 we discuss the overall results and conclude that the adaptive method is the best solution for our problem.

## 2. Test Cases

In this section, we briefly describe the variety of models and observational datasets used as test cases. Our sample selection was based on two criteria:

a) data cubes with low signal-to-noise features such as tails, extra-planar gas and filaments;
b) clean data cubes, i.e. with negligible, or at most minor, artifacts due to calibration and imaging effects.

The consequence of the second criterion is that the filtering results presented in the next sections will be representative for data cubes mainly affected by Gaussian noise.

### *2.1. Models*

We generated several models by taking an existing observation and isolating the detected signal manually. The object (NGC3359), and hence the model, consists of the H I content in a spiral galaxy and a small companion, with an incomplete tidal tail-like structure between them. Gaussian noise has been added with the `GIPSY` (van der Hulst et al., 1992; Vogelaar and Terlouw, 2001) routine RANDOM to produce models with different peak signal-to-noise ratio: 22, 32 and 62, named ModelA, ModelB (see Fig. 1) and ModelC respectively. The signal-to-noise properties of ModelB are the closest to the observational data shown in this paper. Therefore, ModelB will be used as the main reference model.

The data cube size is $6.7 \times 10^5$ voxels. The beam size is $\sim 88'' \times 70''$. The pixel spacing is:

1. $20''$ in Right Ascension (RA), i.e. the data up to $\sim 4$ neighboring pixels are correlated.
2. $20''$ in Declination (Dec), i.e. the data up to $\sim 3$ neighboring pixels are correlated.
3. 8.25 km/s in velocity. The pixels in the velocity direction are not correlated.

These numbers contribute to determining both the optimal width of the filter kernel (see Section 4) and the number of independent voxels, that is $N = 5.6 \times 10^4$.

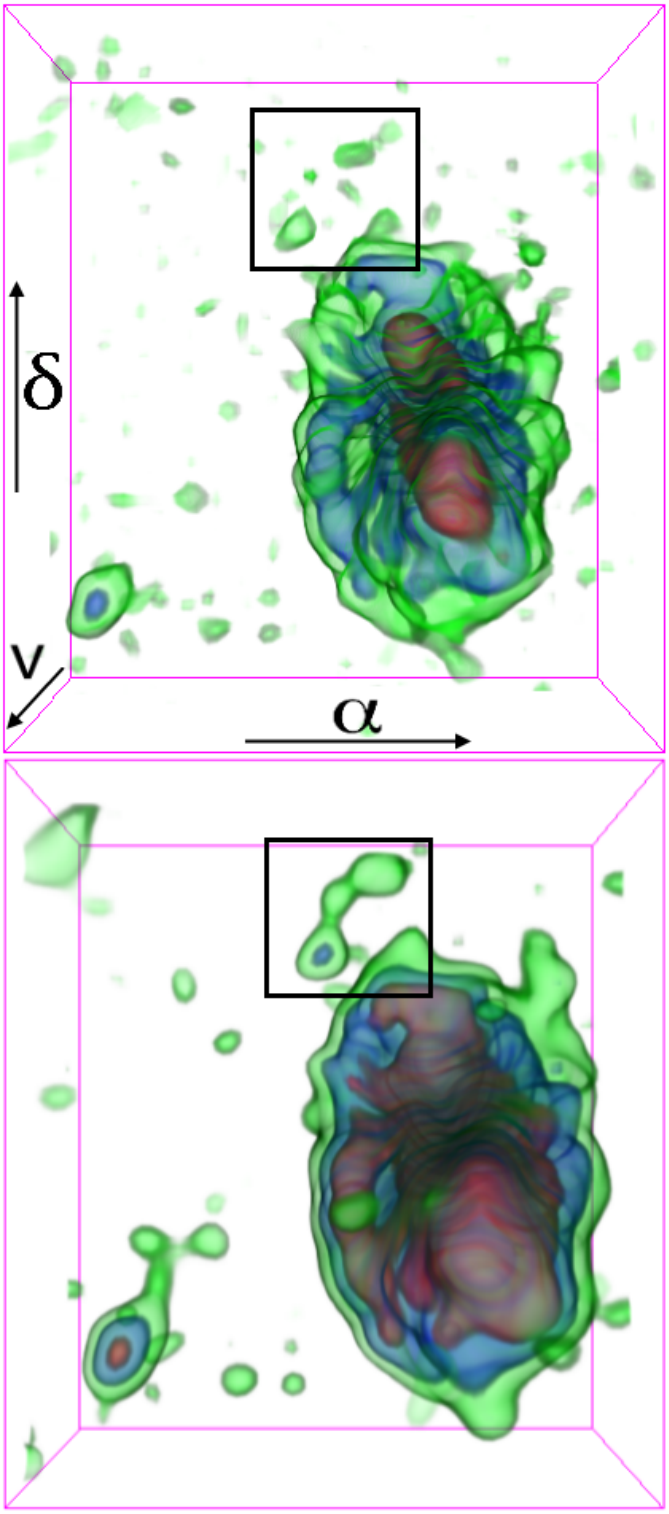


Figure 1: views of modelB. The upper panel shows a volume rendering of the model (information regarding the model is given in Section 2.1) with added Gaussian noise. The bottom panel shows a volume rendering of the smoothed version using an intensity-driven gradient filter with parameters $K = 1.0$, $\tau = 0.0325$, $n = 20$ and $C_{x,y,z} = 5$. The different colors highlight different intensity levels in the data: green, blue and red correspond to 3, 7 and 15 times the *rms* noise respectively. The region of interest, ROI (i.e. the black box), highlights the faint signal, i.e. part of a very faint tail.

### 2.2. NGC4111

NGC4111 is one of the brightest lenticular galaxies in the Ursa Major cluster. The main characteristic of the H I emission of NGC4111 is an extended faint filament between the three sources of the datacube. The orientation and kinematics of this filament suggest that the galaxies were tidally stripped from the outer disks by their nearby companions (Verheijen, 2004).

In Fig. 2, we show a volume rendering of the H I data (Busekool et al., 2016, in prep.) observed with the Very Large Array, VLA, telescope. The size of the data cube is $1.6 \times 10^6$ voxels. The beam size is $\sim 45'' \times 45''$. The pixel spacing is:

1. $15''$ in RA, i.e. the data are correlated up to $\sim 3$ neighboring pixels.
2. $15''$ in Dec, i.e. the data are correlated up to $\sim 3$ neighboring pixels.
3. 5 km/s in velocity. The data are correlated over 2 neighboring pixels because of the use of Hanning smoothing in velocity.

The resulting number of independent voxels is $N = 8.9 \times 10^4$.

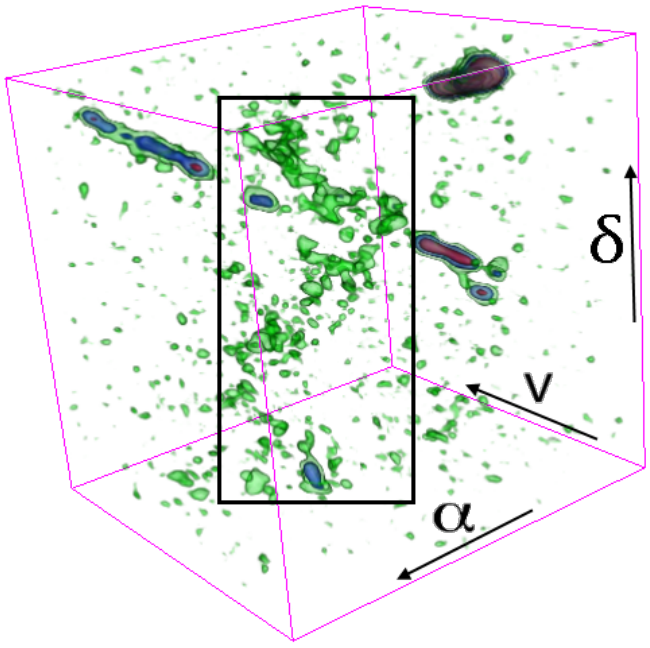


Figure 2: A view of the H I in and around NGC4111 (information regarding the dataset is given in Section 2.2; Busekool et al., 2016, in prep). The different colors highlight different intensity levels in the data: green, blue and red correspond to 3, 7 and 15 times the *rms* noise respectively. The region of interest, ROI (i.e. the black box), highlights the faint signal, i.e. a faint filament between three galaxies.

### 2.3. NGC3379

NGC3379 is an elliptical galaxy in the Leo group. The H I associated with this galaxy is characterized by a very large, extended tail. Part of this tail, such as the wing-shape structure close to the galaxy, is very faint.

In Fig. 3 we show a volume rendering of the H I data observed with the WSRT telescope by Serra et al. (2012). The size of the data cube is $3.6 \times 10^7$ voxels. The beam size is $81'' \times 32''$. The pixel spacing is:

1. $10''$ in RA, i.e. the data are correlated up to $\sim 8$ neighboring pixels.
2. $10''$ in Dec, i.e. the data are correlated up to $\sim 3$ neighboring pixels.
3. 8.25 km/s in velocity. The data are correlated over 2 neighboring pixels because of the use of Hanning smoothing in velocity.

The resulting number of independent voxels is $N = 7.5 \times 10^5$.

### 2.4. WEIN069

The H I data cube of WEIN069 used in this paper is a small sub-cube selected from a large mosaic of 48 WSRT pointings (Ramatsoku et al., 2016), directed towards a region in the sky where a filament of the Perseus-Pisces Supercluster (PPScl) crosses the plane of the Milky Way. The optical counterpart, WEIN069, has been observed by Weinberger et al. (1995).

The data cube is shown in Fig. 4. It contains two sources, WEIN069 and a companion, a tidal tail and a very faint filament that connects the two galaxies. Its size is $7.8 \times 10^5$ voxels. The beam size is $\sim 15'' \times 15''$. The pixel spacing is:

1. $6''$ in RA, i.e. the data are correlated up to $\sim 3$ neighboring pixels.
2. $6''$ in Dec, i.e. the data are correlated up to $\sim 3$ neighboring pixels.
3. 8.25 km/s in velocity. The data are correlated over 2 neighboring pixels because of the use of Hanning smoothing in velocity.

The resulting number of independent voxels is $N = 4.3 \times 10^4$.

## 3. Filtering techniques

In this paper, we focus on interactive filtering of radio data coupled to interactive visualization. The aim is to enhance the manual data inspection, in particular of low signal-to-noise H I structures.

In the next subsections, we list the filtering algorithms used in our analysis in Section 4. The techniques described are aimed to suppress the Gaussian white noise. Moreover, such filters perform well for data with the following characteristics:

i) signal extended over many pixels;
ii) rather small spatial intensity derivatives, i.e. no sharp edges.

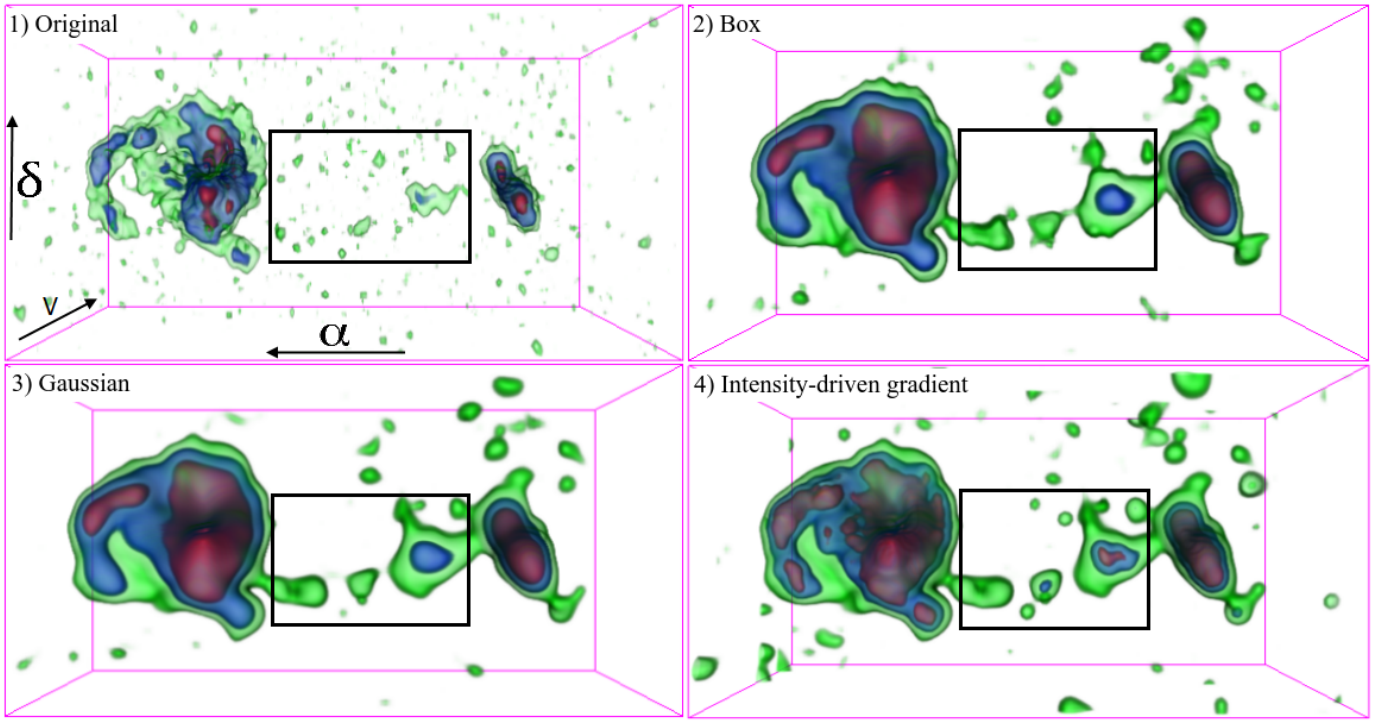


Figure 4: The H I in and around WEIN069 (information regarding the dataset is given in Section 2.4; Weinberger et al., 1995; Ramatsoku et al., 2016). The four panels show: 1) a volume rendering of the original resolution data; 2) the data filtered with a box filter with parameters $N_{x,y,z} = 7$ pixels; 3) the data filtered with a Gaussian filter with parameters $FWHM_{x,y,z} = 5$ pixels; 4) the data filtered with an intensity-driven gradient with parameters $K = 1.5$, $\tau = 0.0325$, $n = 20$ and $C_{x,y,z} = 5$. The different colors highlight different intensity levels in the data: green, blue and red correspond to 3, 7 and 15 times the *rms* noise respectively. The region of interest, ROI (i.e. the black box), highlights the faint signal, i.e. a faint filament between the two companions.

Data of H I in and around galaxies fall into this class. A good example is presented in Fig. 4, one of the data cubes of our sample.

Artifacts generated by effects such as Radio Frequency Interference (RFI), errors in the bandpass calibration or in the continuum subtraction have different statistical properties. Other filter techniques are required to efficiently characterize these artifacts, tailored to their special spatial and spectral signature. In this paper we focus on 'clean' data cubes that are considered free from such artifacts.

For a full review of image processing techniques we refer to Goyal et al. (2012); Buades et al. (2005); Gonzalez and Woods (2002); Weeks (1996).

It is also worthwhile to mention the following automated segmentation methodologies (i.e. automated source mask generation):

1. `SoFiA` (Serra et al., 2015): this pipeline has several tasks for smoothing, source finding and mask optimization. A graphical user interface is also available. Three source-finder algorithms are available: i) a threshold finder; ii) a Smooth and Clip (S-C) finder, which applies thresholding after smoothing the data with a set of user-specified Gaussian kernels and then merges the results; iii) the CNHI finder, which performs a threshold rejection Kuiper test on extracted 1-D spectra. The completeness and reliability of detected sources are evaluated through statistical evaluation of parameters such as the peak flux, total flux, and number of voxels of both positive and negative detections. (Serra et al., 2012).
2. `Duchamp` (Whiting, 2012): this pipeline mainly uses a multi-resolution wavelet transform (specifically the *à trous* algorithm; Starck and Murtagh (1994)) for thresholding the data in the wavelet domain. False detections are rejected using the false discovery rate technique (Hopkins et al., 2002).
3. `MAX – TREE` (Carlinet and Géraud, 2012): this is a tree representation of the data of which the different nodes are classified based on their attributes. These attributes are used to determine the properties of the node (for more information see Teeninga et al. (2015a)). This algorithm has been applied both to interactive visualization (Westenberg et al., 2007) and optical 2-D data (*MT objects*; Teeninga et al., 2015a,b). Preliminary experiments are also ongoing for H I data (*MT source finder*; Moschini et al., 2014; Arnoldus, 2015).

### *3.1. Box filter*

The mean filter (the box filter) simply consists of replacing each pixel value in an image with the mean value of its neighbors, including itself. This has the effect of

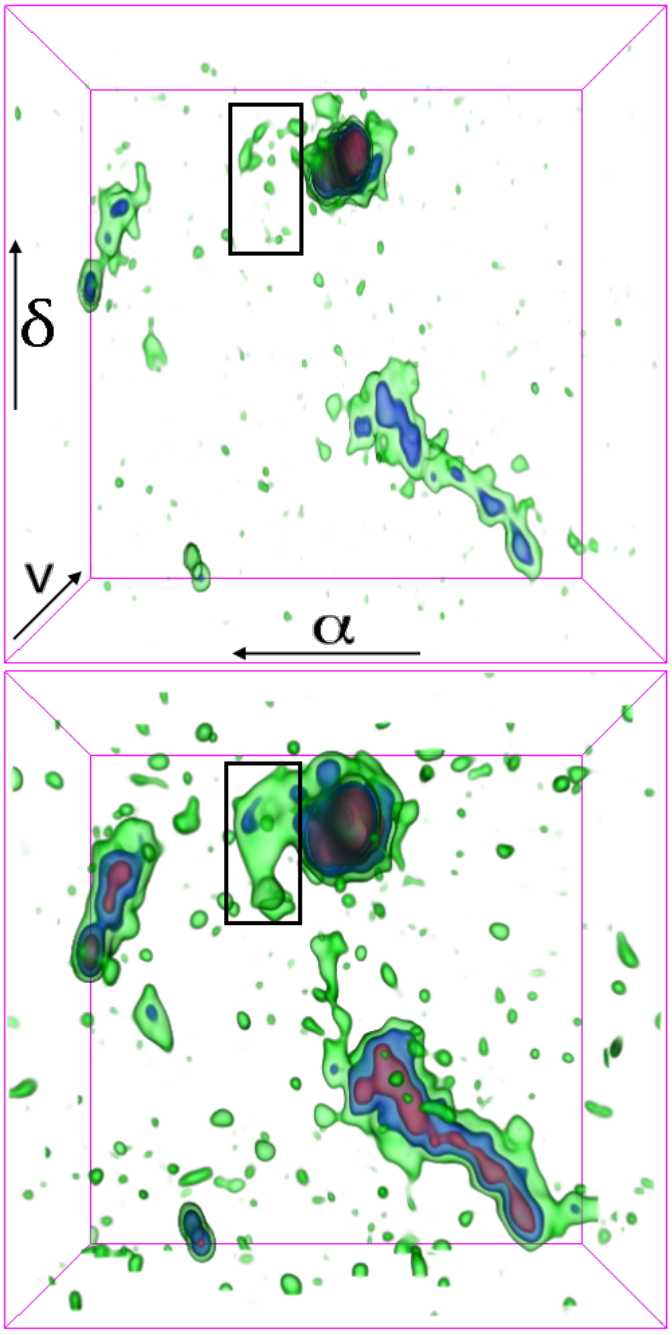


Figure 3: Two views of the H I in and around NGC3379 (information regarding the dataset is given in Section 2.3; Serra et al., 2012). The upper panel is the volume rendering of the original resolution data. The bottom panel shows a volume rendering of the smoothed version using an intensity-driven gradient filter with parameters $K = 1.5$, $\tau = 0.0325$, $n = 20$ and $C_{x,y,z} = 5$. The different colors highlight different intensity levels in the data: green, blue and red correspond to 3, 7 and 15 times the *rms* noise respectively. The region of interest, ROI (i.e. the black box), highlights the faint signal, i.e. a faint wing-shape tidal structure.

eliminating pixel values that are unrepresentative of their surroundings.

The box filter is a convolution filter. Like other convolutions, it is based on a kernel that represents the shape and size of the neighborhood to be sampled when calculating the mean. Box filtering is most commonly used as a simple method for reducing noise in an image (see Fig. 4). However, it has the following drawbacks:

a) a single pixel with a strong artifact, such as RFI, can significantly affect the mean value of all the pixels in its neighborhood;
b) when the filter neighborhood straddles an edge, the filter will blur that edge, leading to a loss of information if the edge is sharp. For H I data this is rarely the case: the effect is visible around the green edges (3 *rms*) of the H I filament (see second panel in Fig. 4). It is a second order effect which only partially degrades the smoothing quality (i.e., the main structure is still visible).

In general, the box filter acts as a low pass filter and, therefore, reduces the spatial intensity derivatives present in the image. The computational complexity of the box filter is $O(N^3)$, where $N$ is the number of voxels.

### *3.2. Gaussian filter*

The Gaussian filter is a 3-D convolution operator that is used to denoise images by smoothing. The kernel is the following Gaussian function:

$$G(x,y,z) = A\exp^{-\left(\frac{(x-x_0)^2}{2\sigma_x^2}+\frac{(y-y_0)^2}{2\sigma_y^2}+\frac{(z-z_0)^2}{2\sigma_z^2}\right)}, \quad (1)$$

where the parameters $\sigma_x$, $\sigma_y$, $\sigma_z$ are related to the full width at half maximum (FWHM) of the peak according to

$$FWHM_i = 2\sqrt{2\ln(2)}\sigma_i,\ i = x,y,z, \quad (2)$$

which determines the degree of smoothing. The 3-D kernel can be also rotated:

$$K(x,y,z) = R_z(\theta_z)\ R_y(\theta_y)\ R_x(\theta_x)\ G(x,y,z), \quad (3)$$

where $R_x$, $R_y$ and $R_z$ are the Euler rotation matrices corresponding to the three Euler angles $\theta_x$, $\theta_y$ and $\theta_z$.

Once a suitable kernel has been calculated, then the Gaussian smoothing can be performed using standard convolution methods. The computational complexity of the Gaussian filter is $O(N^3)$.

When the convolution kernel is isotropic ($\sigma_x = \sigma_y = \sigma_z$), the convolution can be performed much faster since the equation for the 3-D isotropic Gaussian is separable into the three axial components. Thus, the 3-D convolution can be performed with three separate 1-D Gaussian

convolutions. The computational complexity is then lowered to $O(N)$.

The Gaussian filter outputs a weighted average of each pixel's neighborhood, with the average weighted more towards the value of the central pixels. This is in contrast to the box filter's uniformly weighted average. Because of this, a Gaussian provides gentler smoothing and preserves edges better than a similarly sized mean filter (Buades et al., 2005). For H I data, this effect is minor, however it is possible to observe some small differences at the 3 $rms$ level in Fig. 4 (in the 3-D views the faint signal and noise at the 3 $rms$ level are highlighted in green). These discrepancies increase with larger kernels.

In order to increase the local signal-to-noise ratio of the very faint signal, both the box and the Gaussian filter have to use large kernels for the convolution (Buades et al., 2005). The main drawback of these filters is the loss of the spatial information with high signal-to-noise ratio, i.e. the inner region of the galaxy as shown in the second and third panels in Fig. 4. In the next subsection, we will introduce the intensity-driven gradient filter which is designed to deal with this issue by adaptive smoothing depending on the local signal-to-noise ratio and structure in the data.

### *3.3. Intensity-Driven Gradient filter*

The gradient filter (Perona and Malik, 1990) operates on the differences between neighboring pixels, rather than on the pixel values directly. The algorithm, known also as *anisotropic diffusion*, uses a diffusion process described by the following differential equation:

$$\frac{\partial I(x,y,z,t)}{\partial t} = S(x,y,z,t)\,\triangle I(x,y,z,t) + \nabla S(x,y,z,t)\cdot\nabla I(x,y,z,t), \quad (4)$$

where $I$ is the intensity of the pixel and $S$ is the diffusion coefficient. The algorithm was designed for edge detection by choosing:

$$S(x,y,z,t) = \frac{1}{1+\dfrac{|\nabla I(x,y,z,t)|^2}{K^2}}. \quad (5)$$

Instead of having the degree of blurring be dependent on the magnitude of the gradient, it can also be made dependent on other properties, such as the squared image intensity (Perona and Malik, 1990; Arnoldus, 2015):

$$S(x,y,z,t) = \frac{1}{1+\dfrac{I^2(x,y,z,t)}{K^2\ rms^2}}. \quad (6)$$

Substituting equation 6 in equation 4, we obtain a diffusion algorithm which preserves the edges less well, but it adaptively smooths the pixel intensity (i.e. more smoothing for lower signal-to-noise ratio). The second term of equation 4 can be neglected as shown by Perona and Malik (1990) and we use their approach for the discretization of equation 4. The discretized form of this approximation for the $i$-th and $i+1$-th iteration is:

$$I_{i+1} = I_i + \tau\ \frac{C_x\nabla_x I_i + C_y\nabla_y I_i + C_z\nabla_z I_i}{1+\dfrac{I_i^2}{K^2\ rms^2}}, \quad (7)$$

where the algorithm evaluates this expression $n$ times from $i=0$ to $i=n$. $I_i = I_i(x,y,z)$, $\nabla_x I$ indicates the nearest-neighbor differences defined as $[I(x+1,y,z)-I(x,y,z)]+[I(x-1,y,z)-I(x,y,z)]$, $rms$ is the noise level in the data cube and $\tau$, $C_x$, $C_y$, $C_z$ and $K$ are input parameters. The input parameters have the following upper and lower limits: i) $\tau$ ranges in $[0.0025; 0.0625]$; ii) $C_x$, $C_y$ and $C_z$ range in $[0; 10]$; iii) $K$ ranges in $[0.5, 10]$. We define the following default parameters: $K = 1.5$, $\tau = 0.0325$, $n = 20$, $C_x = C_y = C_z = 5$.

The intensity-driven gradient filter is intrinsically adaptive and is therefore a very powerful tool for investigating low signal-to-noise, extended emission such as tails, filaments and extra-planar gas. The fourth panel in Fig. 4 shows an example of gradient smoothing. In the inner part of the galaxy (shown in red at levels above 15 $rms$) the full resolution is conserved remarkably well, while the fainter structure in the outer part shown in green (i.e. the filament at 3 $rms$) has been enhanced at the expense of resolution. A disadvantage of the adaptive smoothing process is that it does not conserve the flux scale. The consequence is that the results can be used for visualization purposes, but not for quantitative analysis. Operations such as calculating column densities, intensity weighted mean velocities, velocity dispersions etc., must be performed on the original data cube or properly convolved versions. The computational complexity of the intensity-driven gradient filter is $O(n\ N)$.

### *3.4. Wavelet filter*

Wavelet transformations are used to obtain a multiresolution representation for analyzing the information content of images. An advantage is that in the wavelet domain it is easier to discriminate the signal from the noise of the image. The decomposition process, mathematically reversible, defines a multiresolution representation (for more information, see Mallat, 1999). In this paper, we restricted ourselves to wavelet transformations using the orthogonal Haar wavelet (Daubechies and Sweldens, 1998) and the biorthogonal Le Gall 5/3 wavelet (known also as Cohen-Daubechies-Feauveau 5/3, CDF 5/3, wavelet; Gall and Tabatabai, 1988).

To obtain a wavelet representation, we used a *wavelet lifting* algorithm (Daubechies and Sweldens, 1998). Wavelet lifting consists in applying low and high pass filters, corresponding to the chosen wavelet, at different resolutions. At each resolution, the low pass filter generates an *approximation band*, $c^l$, and the high pass filter generates

the *detail band*, $d^l$, both of length $N/2^l$ elements, where $l$ is the value of the decomposition level. The approximation bands represent the coarse features in the data, while the detailed bands represent the fine features. The fine features are the differences between the full resolution data and the new coarse version. The detailed bands are used to restore the original data from the coarse resolution.

The wavelet lifting algorithm is performed in 3 steps:

1) Splitting: this step splits a signal into two sets of coefficients, those with even and those with odd index, indicated by $even^l$ and $odd^l$. This is called the *lazy wavelet transform.*
2) Prediction: as the even and odd coefficients are correlated, we can predict one from the other:

$$d^{i,l+1} = odd^{i,l} - P(even^{i,l}), \tag{8}$$

where $i$, is the index for the $i$-*th* array element, and the predict operator, $P$, in the case of the Haar wavelet, is

$$P(even^{i,l}) = even^{i,l}. \tag{9}$$

3) Update: similarly to the prediction step

$$c^{i,l+1} = even^{i,l} + U(odd^{i,l+1}), \tag{10}$$

where the update operator, $U$, for the Haar wavelet, is

$$U(odd^{i,l+1}) = \frac{d^{i,l+1}}{2} \tag{11}$$

An image is then denoised by applying thresholding to the detail bands. Performing wavelet lifting does not require additional memory. In addition, the computational complexity of wavelet lifting is $O(N)$ which makes the algorithm extremely fast.

Wavelet lifting has been widely used as a tool for image denoising in several fields. A practical example of an application of image denoising with wavelet transforms in the case of functional magnetic resonance imaging (fMRI) can be found in Wink and Roerdink (2004).

In Fig. 5, we show filtering results based on the Haar and Le Gall 5/3 wavelets. We pre-smoothed the data with a Gaussian filter with parameters $FWHM_{x,y,z} = 5$, then we decomposed the signal up to the third decomposition level and we finally applied thresholding to the approximation and detail bands. We note that, in general, wavelet denoising algorithms for suppressing Gaussian noise apply thresholding only to the detail bands. However, in the case of H I data, we discovered that it is necessary to threshold both the detail and approximation bands to properly isolate the signal from the noise (the signal is extremely faint). As a result the algorithm is effectively a thresholding filter. The values of the thresholding parameters, $t_{l,wavelet}$, used are: i) $t_{1,Haar} = 0.5$, $t_{2,Haar} = 0.8$ and $t_{3,Haar} = 1.1$ times the $rms$ of the original data cube for the Haar wavelet; ii) $t_{1,LeGAll} = 1$, $t_{2,LeGAll} = 1.4$ and $t_{3,LeGAll} = 1.7$ times the $rms$ of the original data cube for the Le Gall 5/3 wavelet. Throughout this paper the thresholding parameters will always be defined in units of the $rms$ of the original data cube. Comparing the three panels in Fig. 5, one can clearly see that the wavelet filters remove the noise efficiently with minimal loss of the signal.

The algorithms used have, however, some drawbacks. The Haar filter looses resolution at low signal-to-noise ratio due to the averaging of neighborhood pixels. The Le Gall 5/3 filter applies an additional degree of smoothing and generates clear artifacts as shown in Fig. 5.

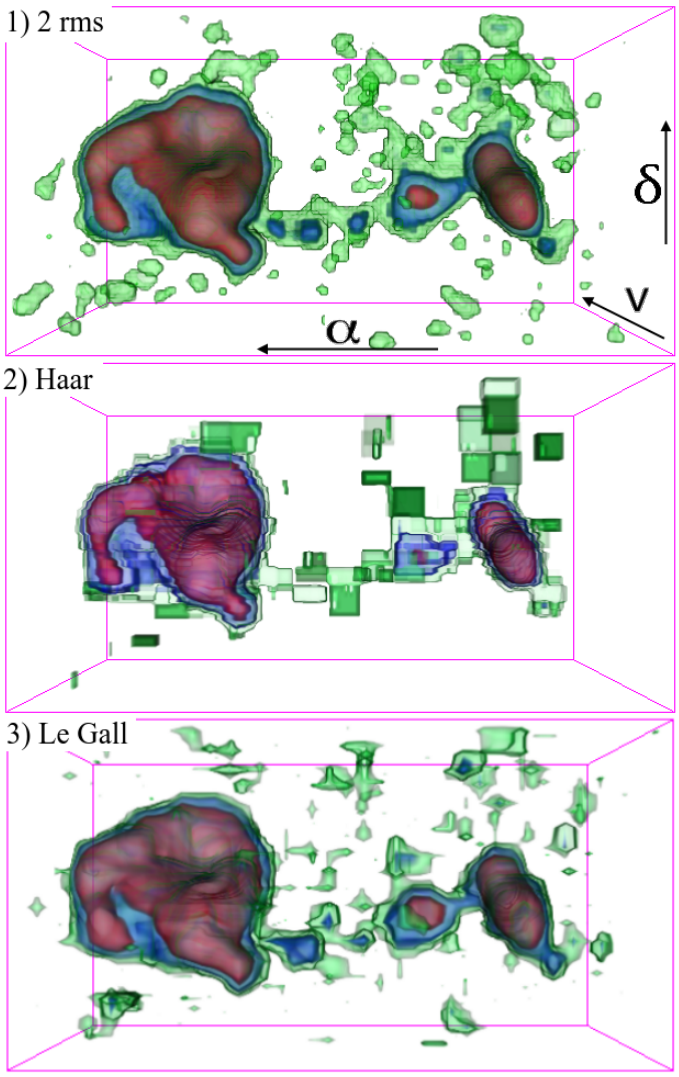


Figure 5: Three volume renderings of WEIN069. In the upper panel, we show the filtered data applying a 2 $rms$ thresholding. In the middle and lower panels the data are filtered with a Haar and Le Gall wavelet thresholding filter, respectively. We performed the decomposition up to level $l = 3$ for both wavelets, then we applied thresholding. In all the cases, we pre-smoothed the data with the same Gaussian filter with parameters $FWHM_{x,y,z} = 5$ pixels. The different colors highlight different intensity levels in the data: green, blue and red correspond to 3, 7 and 15 times the $rms$ noise respectively.

Although the output images obtained by wavelet denoising algorithms are affected by artifacts, wavelet thresholding is very promising when compared to a simple 2 $rms$

thresholding filter. On the other hand, to use wavelet thresholding effectively we encountered the following complications:

1) finding the right multi-level thresholds in the wavelet space is a rather difficult task, which highly depends on the signal-to-noise ratio of the faint signal in the wavelet domain;
2) the choice of the decomposition level, $l$, and of the wavelet highly depends on the spatial and velocity extents of the unknown faint signal (e.g. higher order wavelets may give different results).

A full investigation to determine the optimal wavelet, decomposition level and threshold values for denoising H I data with a wide range of properties will be extremely useful. Flöer and Winkel (2012) provided an analysis and application of wavelet filters for source finding. They demonstrated that separating the wavelet analysis of the spatial dimensions from the velocity dimension increases the filtering quality. However, their study focused on non-resolved galaxies. In the case of well-resolved galaxies the presence of faint and unusual H I structures adds even more complexity to the problem. We will discuss this further in Section 7.

## 4. Optimal filtering parameters

In Section 3, we qualitatively illustrated the filtering results of applying box, Gaussian, intensity-driven gradient, and wavelet lifting algorithms to the WEIN069 data cube (Fig. 4). In this section, we compare quantitatively the box, Gaussian and intensity-driven gradient filtering output, for the full sample defined in Section 2. In order to quantify the smoothing quality, we define a diagnostic parameter:

$$F = \frac{|S_{o,f}|}{|S_{o,o}|}, \tag{12}$$

where

$$\begin{aligned} S_{o,f} &= \sum_{(x,y,z)} I_o(x,y,z)\; M_f(x,y,z) \\ S_{o,o} &= \sum_{(x,y,z)} I_o(x,y,z)\; M_o(x,y,z), \end{aligned} \tag{13}$$

$$M_i(x,y,z) = \begin{cases} 1 & \text{if } I_i(x,y,z) \geqslant 3\, rms_i,\ i = o, f \\ 0 & \text{if } I_i(x,y,z) < 3\, rms_i. \end{cases} \tag{14}$$

In the previous equations, $rms_i$ is the root mean square (i.e. noise level), $I_i(x,y,z)$ is the intensity of the pixel at coordinates $(x,y,z)$, the index $o$ refers to the original data cube and $f$ to the filtered one. The coordinates $(x,y,z)$ range in a ROI sub-cube of a faint signal as shown (with a black box) in Figs. 1, 2, 3 and 4. Moreover, the values of the sums $S_{o,o}$ and $S_{o,f}$, in equation 13, are always calculated on the pixel intensity values of the original data cube. Therefore, it represents a measurement independent of the filtering technique used.

The $F$ parameter can range between $[0, M]$ where $M$ is an unknown upper limit (see Section 5). The parameter has a different meaning depending on its range:

i) $F \in [0,1]$: the smoothing has washed out the faint signal. This can easily happen using box or Gaussian kernels that are too large.
ii) $F \in [1, M]$: the faint signal has been enhanced and the number of voxels in the mask $M_f$ is generally larger than in $M_o$. The $F$-value is correlated with the smoothing quality. For high values of $F$, the filtered data cube has more signal raised over its 3 $rms$ noise level.

The error, $\sigma_F$, is propagated as:

$$\begin{aligned} \sigma_F &= \sqrt{\left(\frac{\partial F}{\partial S_{o,f}}\right)^2 \sigma^2_{S_{o,f}} + \left(\frac{\partial F}{\partial S_{o,o}}\right)^2 \sigma^2_{S_{o,o}}} = \\ &\sqrt{\frac{\sigma^2_{S_{o,f}}}{S^2_{o,o}} + \frac{S^2_{o,f} * \sigma^2_{S_{o,o}}}{S^4_{o,o}}}, \end{aligned} \tag{15}$$

where $S_{o,o}$ and $S_{o,f}$ are affected by an error due to the Gaussian noise background equal to

$$\sigma_{S_{o,i}} = \sqrt{N_i}\, rms_i. \tag{16}$$

In the last equation $i$ is an index which is either $f$ or $o$, $N_i$ is the number of independent voxels in the mask $M_i$ and we assumed the $rms_i$ to be constant in the full data cube.

We report the values of the $F$ parameter ($F$-values) in Table 1, in which the best runs and their parameters are reported for each data cube and filter. The results shown in the table are due to a fine-tuning process of the parameter space based both on visual inspection of the data and evaluation of the $F$-values. The specific input parameter space for each algorithm is:

1) box filter: $N_j = 1, 3, 5$ for the Models; $N_j = 5, 7, 9$ for WEIN069, NGC4111 and NGC3379;
2) Gaussian filter: $FWHM_j = 1, 3, 5$ for the Models; $FWHM_j = 3, 5, 7$ for WEIN069, NGC4111 and NGC3379;
3) wavelet filter: $l = 1, 2, 3$;
   $t_{1,Haar} = 0.1, 0.3, 0.5, 0.7, 0.9$,
   $t_{2,Haar} = 0.4, 0.6, 0.8, 1.0, 1.2$,
   $t_{3,Haar} = 0.7, 0.9, 1.1, 1.3, 1.5$,
   $t_{1,LeGall} = 0.6, 0.85, 1.1, 1.35, 1.6$,
   $t_{2,LeGall} = 0.9, 1.15, 1.4, 1.65, 1.9$,
   $t_{3,LeGall} = 1.2, 1.45, 1.7, 1.95, 2.2$;
   we also pre-smoothed the data with a Gaussian filter with parameters $FWHM_{x,y,z} = 3$ for the models; $FWHM_{x,y,z} = 5$ for WEIN069, NGC3379 and NGC4111;
4) intensity-driven gradient filter: $K = 0.5, 1, 1.5, 2$; $n = 20, 30$; $\tau = 0.0325, 0.0625$; $C_j = 4, 5, 6$;

| data cube | Filter | Best input parameters | $F$ |
|---|---|---|---|
| ModelA | 1 | $N_x = 3$ ; $N_y = 3$ ; $N_z = 1$ | $1.1996 \pm 0.0307$ |
| | 2 | $FWHM_x = 3$ ; $FWHM_y = 3$ ; $FWHM_z = 1$ | $1.1946 \pm 0.0286$ |
| | 3 | Haar wavelet ; $l = 2$ ; $t_1 = 0.7$ ; $t_2 = 1.0$ | $3.8251 \pm 0.0778$ |
| | 4 | $K = 1$ ; $\tau = 0.0325$ ; $n = 20$ ; $C_x = 4$ ; $C_y = 5$ ; $C_z = 4$ | $2.1787 \pm 0.0463$ |
| ModelB | 1 | $N_x = 3$ ; $N_y = 5$ ; $N_z = 3$ | $1.2975 \pm 0.0196$ |
| | 2 | $FWHM_x = 3$ ; $FWHM_y = 5$ ; $FWHM_z = 3$ | $1.5959 \pm 0.0228$ |
| | 3 | Haar wavelet ; $l = 2$ ; $t_1 = 0.7$ ; $t_2 = 0.8$ | $2.8515 \pm 0.0394$ |
| | 4 | $K = 1$ ; $\tau = 0.0325$ ; $n = 20$ ; $C_x = 5$ ; $C_y = 5$ ; $C_z = 4$ | $2.3280 \pm 0.0339$ |
| ModelC | 1 | $N_x = 5$ ; $N_y = 5$ ; $N_z = 1$ | $1.0928 \pm 0.0078$ |
| | 2 | $FWHM_x = 5$ ; $FWHM_y = 5$ ; $FWHM_z = 1$ | $1.2312 \pm 0.0082$ |
| | 3 | Haar wavelet ; $l = 2$ ; $t_1 = 0.9$ ; $t_2 = 0.6$ | $1.1062 \pm 0.0072$ |
| | 4 | $K = 1$ ; $\tau = 0.0325$ ; $n = 20$ ; $C_x = 6$ ; $C_y = 6$ ; $C_z = 4$ | $1.6407 \pm 0.0109$ |
| WEIN069 | 1 | $N_x = 9$ ; $N_y = 7$ ; $N_z = 7$ | $1.9967 \pm 0.0068$ |
| | 2 | $FWHM_x = 7$ ; $FWHM_y = 5$ ; $FWHM_z = 5$ | $2.2576 \pm 0.0076$ |
| | 3 | Le Gall wavelet ; $l = 3$ ; $t_1 = 0.6$ ; $t_2 = 1.9$ ; $t_3 = 1.45$ | $2.8999 \pm 0.0096$ |
| | 4 | $K = 1.5$ ; $\tau = 0.0325$ ; $n = 20$ ; $C_x = 6$ ; $C_y = 5$ ; $C_z = 4$ | $2.3392 \pm 0.0081$ |
| NGC4111 | 1 | $N_x = N_y = N_z = 9$ | $3.0789 \pm 0.0032$ |
| | 2 | $FWHM_x = FWHM_y = FWHM_z = 7$ | $3.3057 \pm 0.0034$ |
| | 3 | Le Gall wavelet ; $l = 3$ ; $t_1 = 1.1$ ; $t_2 = 0.9$ ; $t_3 = 1.2$ | $3.7505 \pm 0.0036$ |
| | 4 | $K = 2$ ; $\tau = 0.0325$ ; $n = 30$ ; $C_x = 5$ ; $C_y = 6$ ; $C_z = 5$ | $2.9665 \pm 0.0031$ |
| NGC3379 | 1 | $N_x = N_y = N_z = 9$ | $5.6655 \pm 0.0078$ |
| | 2 | $FWHM_x = FWHM_y = FWHM_z = 7$ | $5.9252 \pm 0.0081$ |
| | 3 | Le Gall wavelet ; $l = 3$ ; $t_1 = 0.6$ ; $t_2 = 1.15$ ; $t_3 = 1.2$ | $6.3993 \pm 0.0233$ |
| | 4 | $K = 2$ ; $\tau = 0.0325$ ; $n = 30$ ; $C_x = C_y = C_z = 6$ | $5.2800 \pm 0.0072$ |

Table 1: Best runs are reported. We performed the selection evaluating the $F$-values and confirming it by visual inspection. The filter index entries are respectively: 1) box; 2) Gaussian; 3) wavelet lifting thresholding (with Gaussian pre-smoothing); 4) intensity-driven gradient;; The parameters $N$ and $FWHM$ are defined in pixel units. The parameters $t_{l,wavelet}$ are defined in units of $rms$ noise level of the original data cube. The parameters $l$, $K$, $\tau$, $n$ and $C$ are dimensionless.

where $j = x, y, z$. Note that a detailed tuning parameter search can be performed iteratively at higher resolutions (Bergner et al., 2013). However, the input parameter sample used is accurate enough for finding optimal $F$-values and, therefore, for judging which are the best input parameters. This has been checked by performing the analysis also with a higher resolution sampling of the input parameters.

Moreover, in our parameter space investigation, we chose to set the rotation parameters for the Gaussian filter, $\theta_i$, to zero to reduce the large input parameters space. This does not introduce a substantial bias in our investigation because the dependencies of the results on the rotation parameters are negligible. In fact, for our sample only filtering results for WEIN069 show a dependence of the $F$-parameter on the Euler rotation angles. In the other cases the faint signal is mainly oriented along one of the primary axes, e.g. NGC3379, or it has a more complex morphology such as the S-shaped filament in NGC4111 or arc-shaped tail in the models. As example, in Table 2, we report the $F$-values of filtering WEIN069 with a rotated Gaussian kernel. The results show that a particular rotation, run III, $\theta_y = 340°$, increases the $F$ parameter by a factor of 7.5%, while in run II, $\theta_z = 340°$, it is smaller by a factor of 9.2%. This is expected, in fact, since most of the faint signal is aligned along a diagonal axis, corresponding to the x-axis rotated by 340°with respect to the y-axis. Therefore, in run III the kernel is aligned to the faint signal, while in run II it is perpendicular to it.

| Run | $\theta_x(°)$ | $\theta_y(°)$ | $\theta_z(°)$ | $F$ | $\sigma_F$ |
|---|---|---|---|---|---|
| I | 0 | 0 | 0 | 2.0701 | 0.0071 |
| II | 0 | 0 | 340 | 1.9776 | 0.0068 |
| III | 0 | 340 | 0 | 2.1447 | 0.0073 |

Table 2: The $F$-values applying to WEIN069 a Gaussian filter with parameters $FWHM_x = 7$ pixels, $FWHM_{y,z} = 3$ pixels, $\theta_x$, $\theta_y$ and $\theta_z$.

For the wavelet filters, we performed a pre-smoothing step with a Gaussian filter. This was necessary for increasing the signal-to-noise ratio and providing the optimal results shown in this paper using as maximum decomposition level $l = 3$. We experienced that, in the case of H I data, performing a Haar or Le Gall wavelet analysis beyond the third decomposition level gives rise to many artifacts.

In the next section, we will show detailed tests of the $F$-parameter to establish that this parameter is a reliable estimator of the quality of the filtering results. In Section 6 we will present performance benchmarks of our parallel implementation of the filtering algorithms, and show

that parallelization is necessary to satisfy the interactivity requirement defined in Section 1.

## 5. Noise consideration

In this section, we further investigate the $F$-parameter defined in equation 12 and its relations with the signal and noise. In fact, the sum over a pure Gaussian noisy sub-cube is affected by a statistical error equal to $\sqrt{N}$ $rms$ (i.e. the average differs from the zero value). Moreover, applying the mask calculated in the smoothed data cube, $M_f$, to the original data adds further complications: inside the mask there will be a part of the faint signal (e.g. the peak in the histogram of the middle panel in Fig. 6) and partially noise (e.g. the left wing of the same histogram).

In the lower panel in Fig. 6 we show a plot of the $F$-values calculated from masks obtained by spanning the thresholding value of the mask $M_f$ from zero to 4.5 $rms$. We performed the calculation both on the sub-cube containing the faint signal (ROI defined in Fig. 1) and three different sub-cubes, of the same dimension as the ROI, in which there is only noise. In the case of the ROI sub-cube, $F$ increases with increased threshold. Vice versa, for the noise sub-cubes, $F$ decreases with increasing threshold and its value is $\sim 0$ above 2.5 $rms$. Note that the threshold used in Section 4 for the masks $M_i$ is 3 $rms$.

We define also the following parameter:

$$F_M = \frac{|S_{m,f}|}{|S_m|}, \tag{17}$$

where

$$\begin{aligned} S_{m,f} &= \sum_{(x,y,z)} I_m(x,y,z)\; M_f(x,y,z) \\ S_m &= \sum_{(x,y,z)} I_m(x,y,z). \end{aligned} \tag{18}$$

In these equations, the index $m$ indicates the ModelB cube without the Gaussian artificial noise. $S_m$ is the integrated flux over the full ROI sub-cube, therefore $F_M$ is the percentage of recovered signal in the mask $M_f$ and it ranges in [0;1]. $M_f$ is defined in equation 14.

In table 3 we report the $F_M$-values obtained by performing the intensity-driven gradient filter on the ModelB data cube. The table shows that an increase of the parameter $F$ corresponds to an increase of the parameter $F_M$, i.e. more signal has been recovered in the smoothing process. This is also supported by visual inspection of the filtered data cubes. In Fig. 7, the faint signal is clearly enhanced for higher values of $F$ and $F_M$.

We performed the same analysis for ModelA and ModelC, with similar results as the analysis performed on ModelB.

We conclude that the $F$-values are reliable and the noise effects on the $F$-values, calculated at the 3 $rms$ noise level, are minor or negligible.

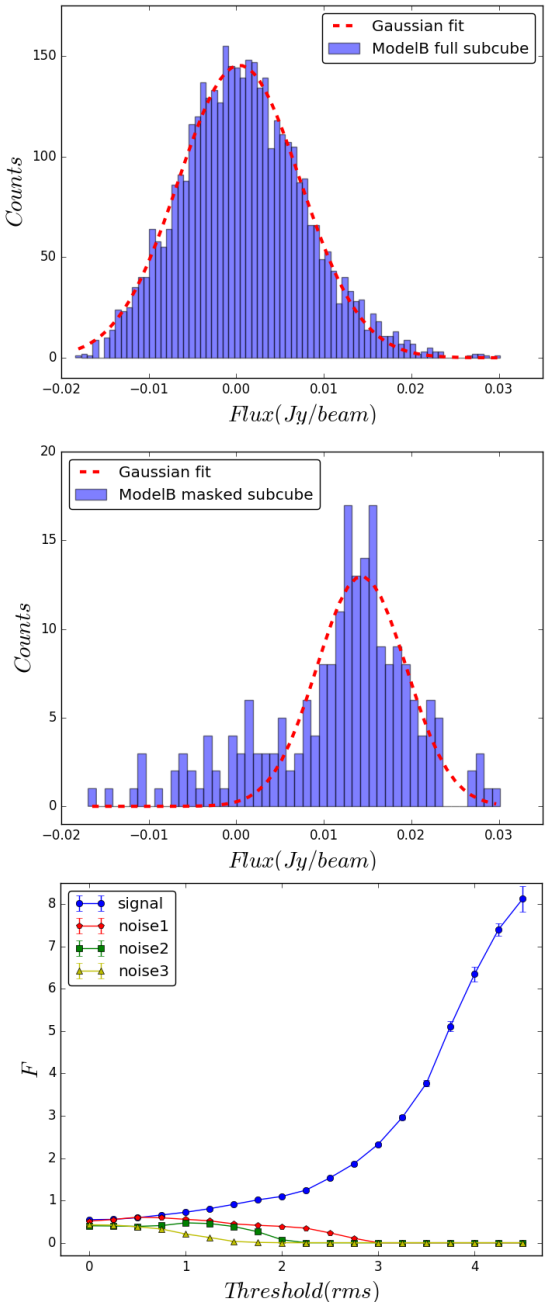


Figure 6: The analysis of the histogram of the pixel intensity distribution and $F$-parameter for ModelB. The upper panel shows the histogram of a sub-cube of ModelB. The sub-cube selection is the ROI, the faint signal, defined in Fig. 1. The red curve is a Gaussian fit over the histogram. The fitted parameters are: $\mu = 3.5\ 10^{-4} \pm 1.0\ 10^{-8}$; $\sigma = 7.0\ 10^{-3} \pm 1.0\ 10^{-8}$; bins = 75. In the middle panel, the histogram applying the mask $M_f$ from run 5 (defined in Table 3) on the ROI is shown. The output parameters of the fitting are: $\mu = 1.4\ 10^{-2} \pm 1.6\ 10^{-7}$; $\sigma = 5.1\ 10^{-3} \pm 1.6; 10^{-7}$; bins = 50. The lower panel is a plot of the $F$-values calculated from masks obtained by spanning the thresholding values of the mask $M_f$ from zero to 4.5 $rms$. The blue line corresponds to the $F$-values calculated on the ROI sub-cube. The red, green and yellow lines correspond to the $F$-values calculated on three different sub-cube, of the same dimension of the ROI, in which there is only noise.

| Run | $K$ | $\tau$ | $n$ | $F$ | $F_M$ |
|---|---|---|---|---|---|
| 1 | 0.5 | 0.0325 | 20 | $2.128 \pm 0.032$ | $0.237 \pm 0.004$ |
| 2 | 0.5 | 0.0625 | 20 | $1.322 \pm 0.021$ | $0.149 \pm 0.003$ |
| 3 | 0.5 | 0.0325 | 30 | $2.193 \pm 0.033$ | $0.263 \pm 0.003$ |
| 4 | 0.5 | 0.0625 | 30 | $1.053 \pm 0.018$ | $0.129 \pm 0.003$ |
| 5 | 1 | 0.0325 | 20 | $2.328 \pm 0.034$ | $0.364 \pm 0.003$ |
| 6 | 1 | 0.0325 | 30 | $2.148 \pm 0.031$ | $0.334 \pm 0.003$ |
| 7 | 1.5 | 0.0325 | 20 | $1.722 \pm 0.025$ | $0.259 \pm 0.002$ |
| 8 | 1.5 | 0.0325 | 30 | $0.633 \pm 0.010$ | $0.071 \pm 0.001$ |
| 9 | 2 | 0.0325 | 20 | $0.701 \pm 0.011$ | $0.079 \pm 0.001$ |
| 10 | 2 | 0.0325 | 30 | $0.263 \pm 0.006$ | $0.024 \pm 0.001$ |

Table 3: The $F$ and $F_M$-values for applying an intensity-driven gradient filter with parameters $K$, $\tau$, $n$, $C_{x,y,z} = 5$ to ModelB.

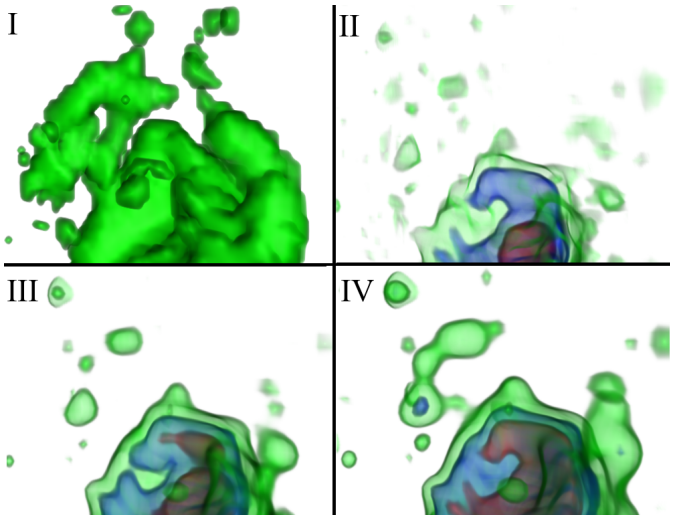


Figure 7: In the four views, we look at a zoom of the ROI defined in Fig. 1. The four panels present the same visualization of four different data cubes: I) the ModelB without the noise; II) ModelB; III) the filtering output obtained by run 4 ($F = 1.053$; see Table 3); IV) the output from run 5 ($F = 2.328$; see Table 3). The different colors highlight different intensity levels in the data: green, blue and red correspond to 3, 7 and 15 times the *rms* noise respectively. The model in the first panel has a *rms* value equal to zero. Therefore, we show only the green level.

## 6. Performance

In this section, we provide measurements of the performance of the codes[1] used in this paper. We performed the benchmark on a Linux laptop (Ubuntu 15.10) equipped with:

- an Intel i7 2.60 GHz CPU,
- 16 GB of DDR3 1.6 GHz *random access memory*, RAM,
- an Intel HD Graphics 4600 *graphics processing unit*, GPU, (it can use up to 1.7 GB of the RAM),
- an NVIDIA GeForce GTX860M GPU (with 2 GB of *dynamic random-access memory*, DRAM).

[1]The codes are publicly available at `https://github.com/Punzo/SlicerAstro`

We define the speedup, $S$, as

$$S(N) = T_1(N)/T_p(N), \tag{19}$$

where $T_1$ is the execution time exploiting only one CPU core, $T_p$ is execution time of the parallelized code and $N$ the number of voxels. The codes are parallelized both on CPU (OpenMP) and GPU (OpenGL). In the case of the GPU implementation, the I/O times (i.e. times for sending the data to the GPU and to getting the results back) are included in the term $T_p$.

We report the speedup results in Fig. 8, using the following values for the input parameters of the filters:

1) isotropic box: $N_x = N_y = N_z = 3$ pixels;
2) anisotropic box: $N_x = 3$ pixels and $N_y = N_z = 5$ pixels;
3) isotropic Gaussian: $FWHM_x = FWHM_y = FWHM_z = 3$ pixels and $\theta_x = \theta_y = \theta_z = 0°$
4) anisotropic Gaussian: $FWHM_x = 2$ pixels, $FWHM_y = FWHM_z = 3$ pixels and $\theta_x = \theta_y = \theta_z = 0°$;
5) intensity-driven gradient: $K = 1.5$, $\tau = 0.0325$, $n = 20$ and $C_x = C_y = C_z = 5$.

A number of conclusions can be drawn from Fig. 8. First, the values of the speedup $S$ for the CPU (8 cores) implementation for the various filters at different $N$ are $\lesssim 4$. Therefore, the execution time for both the anisotropic box and Gaussian filter at large $N$ ($\sim 10^8$) is rather long: 1 min and 10 min, respectively. The isotropic Gaussian filter at large $N$ takes 14.5 s using 8 cores, while the execution time for the isotropic box filter is 2 s, and 56 s for the intensity-driven gradient filter. We also compared our optimized CPU implementation of the isotropic Gaussian filter with the one provided by the Insight Segmentation and Registration Toolkit (`ITK`; Yoo et al., 2002). Our implementation, using the same number of CPU cores (i.e. 8), showed a speedup by a factor of 3 over the `ITK` version.

Secondly, very large values of $S$ are found for the GPU implementation of the anisotropic box and Gaussian filters. For example, in the case of anisotropic Gaussian filtering of NGC2841 ($N = 1.4 \times 10^8$, i.e. 529 MBytes; Walter et al., 2008), the execution time improves from 35 min, using one CPU core, to 3.5 s exploiting the GTX860M.

Thirdly, the values of $S$ for the GPU implementation are smaller for the isotropic box and Gaussian filters than for their anisotropic counterparts. The GPU execution time for the isotropic Gaussian filter with $N = 1.4 \times 10^8$ is 1.8 s and therefore a factor of 2 smaller than for the anisotropic Gaussian filter.

On the other hand the GPU execution time for the intensity-driven gradient filter with $N = 1.4 \times 10^8$ is 3.1 s, as compared to 4 min with a single CPU core.

When examining the behavior in relation to the number of voxels the following conclusions can be drawn. For a data cube with a small number of voxels ($N \sim 10^6$) the $S$ values of the GPU implementation for the isotropic

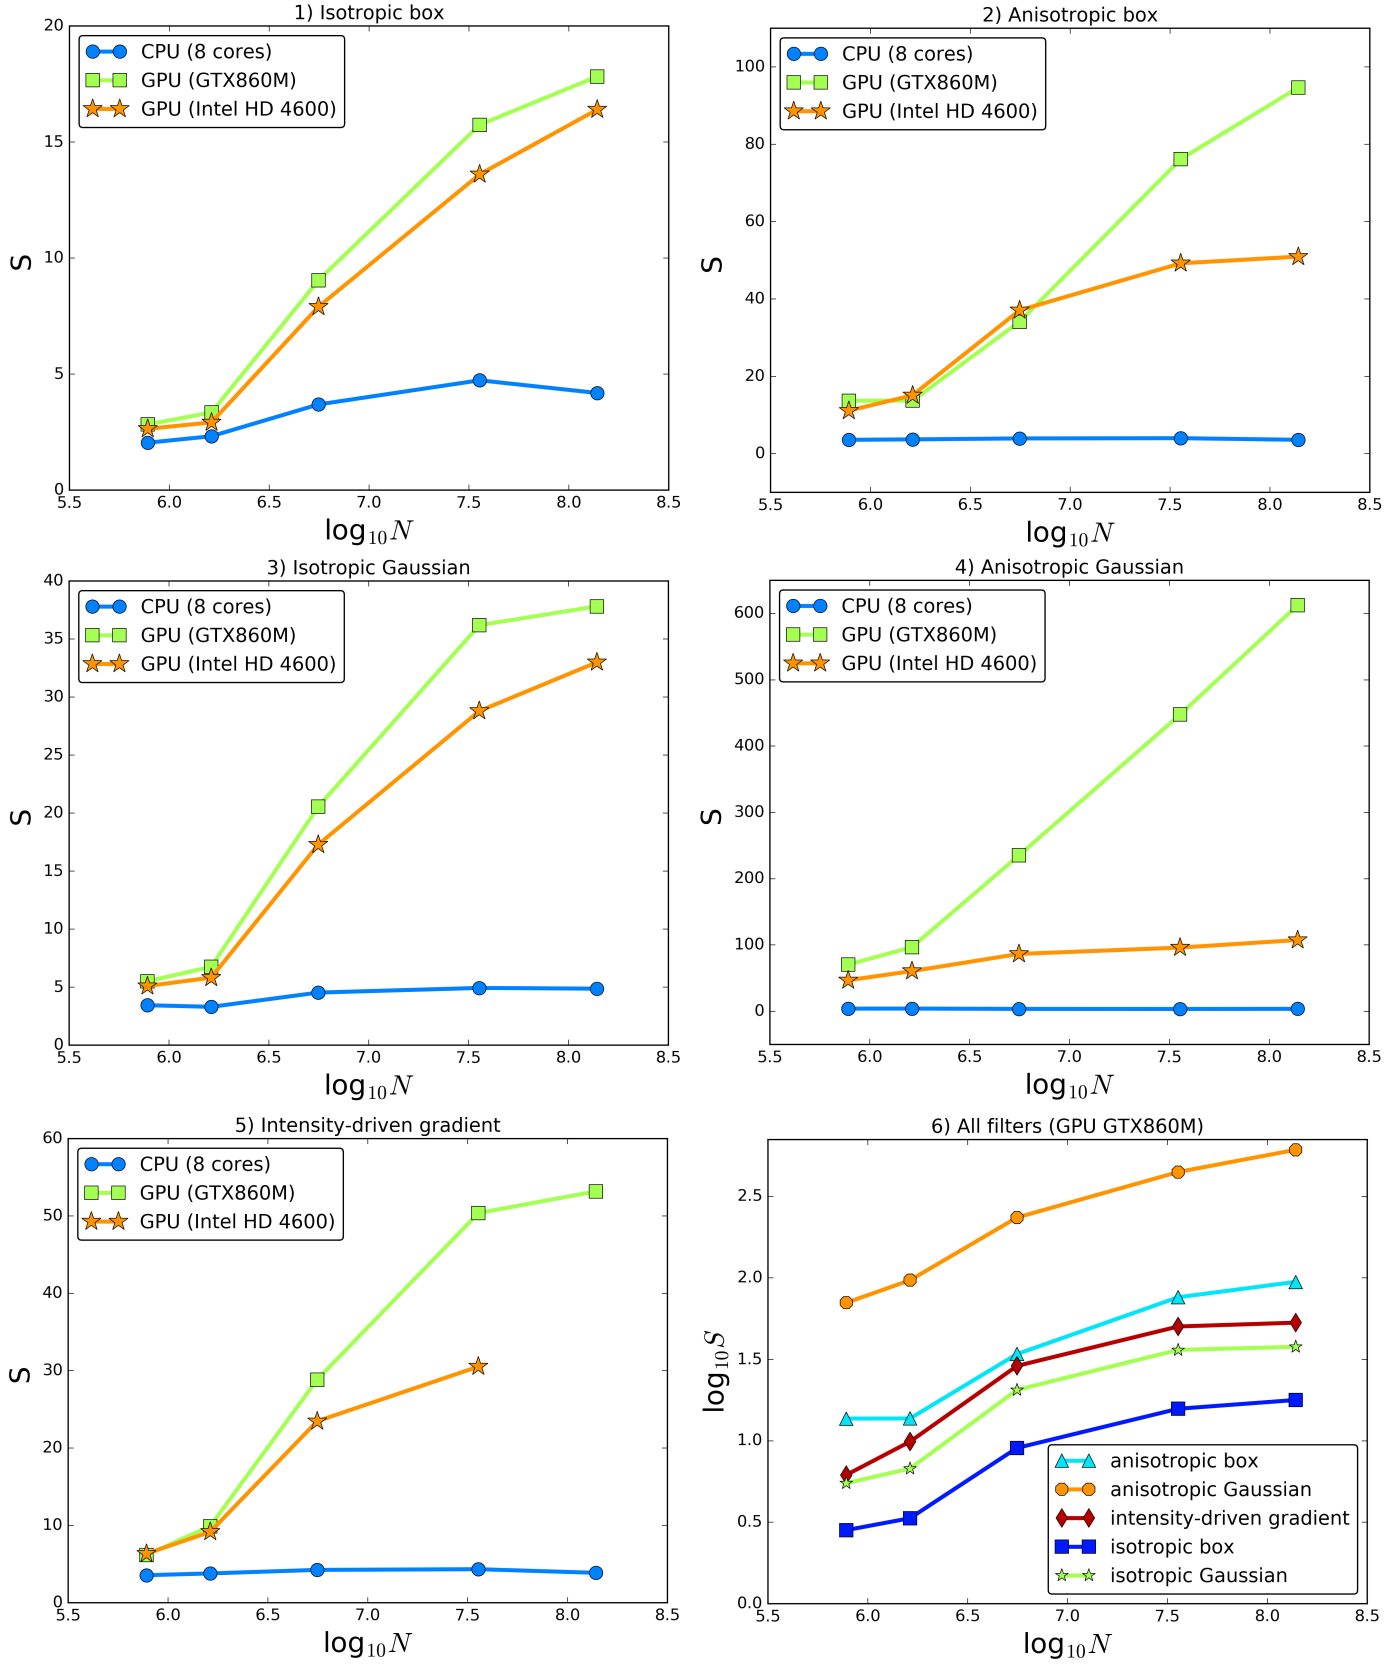


Figure 8: The values of speedup of the parallelization of the various filter algorithms are shown: 1) upper-left panel, isotropic box; 2) upper-right panel, anisotropic box; 3) middle-left panel, isotropic Gaussian; 4) middle-right panel, anisotropic Gaussian; 5) bottom-left panel, intensity-driven gradient; 6) bottom-right panel, comparison of the GPU (GTX860M) implementation of all filters. The values of the speedup $S$ are calculated using equation 19. $N$ is the number of voxels. The values of the input parameters are defined in section 6.

box, isotropic Gaussian and intensity-driven gradient filters are close to the 8 CPU cores performances. This is to be expected as for small $N$ it is not possible to fully load the GPU and properly exploit all the cores. However, up to $5 \times 10^6$ voxels, all filters, when using the GPU, reach the kind of performance that allows interactive work (maximum execution time, exploiting the GTX860M, is $T_p < 0.3s$).

Finally, wavelet lifting is a very fast algorithm: the maximum execution time (using the Haar wavelet and a value of $l = 3$), exploiting one CPU core, for filtering a data cube with up to $10^8$ voxels is $T_p < 5.1\ s$. Therefore, we did not implement a GPU version. Moreover, the implementation of such parallelization is rather challenging mainly because of the memory handling on the GPU. A `CUDA` implementation was developed by Laan et al. (2011) giving a speedup of $\sim 10$ with respect to their optimized CPU implementation. This is a large improvement respect to previous works (e.g. Wong et al. (2007); Tenllado et al. (2008)).

Defining $\eta = 4\ N$ Bytes as the RAM usage for a given data cube, the memory requirements for each of the filter codes are:

A) CPU implementations of the box, Gaussian and intensity-driven gradient filters: one permanent $\eta$ on the RAM for storing the final results and one temporary $\eta$ for storing partial run-time results, so a total memory requirement of $2\eta$ RAM;
B) CPU implementation of the wavelet filters: one permanent $\eta$ on the RAM for storing the final results;
C) GPU implementation of the box, Gaussian and intensity-driven gradient filters: one permanent $\eta$ on the RAM, one temporary $\eta$ on the RAM and two temporary $\eta$ on the DRAM, so a total memory requirement of $2\eta$ RAM and $2\eta$ DRAM;

In summary, a machine with 16 GB of memory can easily accommodate a $\sim$ 4 GB dataset when using the box, Gaussian or intensity-driven gradient filter (in case of the GPU implementation at least 8 GB of DRAM are needed).

For the GPU implementation, we chose the shader paradigm (OpenGL), over other computational scientific SDK (CUDA or OpenCL), for its compatibility with all the GPU vendors. Moreover, OpenGL is present in any operating system, which simplifies the distribution of the software. The drawback is that the computations performed with OpenGL have relatively less precision. For H I data this is not an issue: the scalar range of the pixel intensities is relatively small and float precision is sufficient for the calculations required by the algorithms. In fact, the differences between the CPU and GPU filtered data cubes are unnoticeable: in Table 4, we compare the GPU methods smoothing quality to the CPU ones calculating the $F$-values, and the differences between the two implementations are less than 5%.

| Filter | Hardware | $F$ |
|---|---|---|
| 1 | CPU | 1.7534 ± 0.0061 |
| | GPU | 1.7217 ± 0.0060 |
| 2 | CPU | 1.8591 ± 0.0064 |
| | GPU | 1.8182 ± 0.0063 |
| 3 | CPU | 1.6953 ± 0.0059 |
| | GPU | 1.5386 ± 0.0054 |
| 4 | CPU | 1.8848 ± 0.0065 |
| | GPU | 1.7760 ± 0.0062 |
| 5 | CPU | 2.3416 ± 0.0083 |
| | GPU | 2.2704 ± 0.0079 |

Table 4: The $F$-values relative to both the CPU and GPU filtering implementation of the filters applied to WEIN069. The filter index entries are respectively: 1) isotropic box; 2) anisotropic box; 3) isotropic Gaussian; 4) anisotropic Gaussian; 5) intensity-driven gradient. The values of the input parameters are defined in section 6.

In the next section, we will summarize and discuss the results presented in the previous sections focusing on their applicability to visualization.

## 7. Discussion and conclusions

Future blind surveys of H I will deliver a large variety of data in terms both of the number of galaxies and additional complex features such as tails, extra-planar gas and filaments. These faint structures can be found in nearby medium/high resolved galaxies (e.g. Model and WEIN069 data cube) and groups of non-resolved galaxies (e.g. NGC-3379 and NGC4111). They have a very low signal-to-noise ratio of $\sim$ 1, but are extended over many pixels. Efficiently separating such signals from the noise is not straightforward (visual examples are shown in sections 2 and 3). Moreover, in the case of APERTIF and ASKAP, it is estimated that tens of such sub-cubes will be collected weekly (Duffy et al., 2012). This is a large volume of data, and a coupling between the filtering algorithms shown in this paper and 3-D visualization can enhance the inspection process of large numbers of galaxies and masks provided by source finder algorithms.

In Section 3, we reviewed state-of-the-art filtering algorithms. We qualitatively illustrated the filtering results using several methods. We then performed a visual inspection of the filtering results, followed by a systematic quantitative analysis of the algorithms in Section 4.

First, we extensively investigated the parameter space of the input parameters (i.e. the extension and shape of the kernels) of the box and Gaussian filters by applying them to several test data cubes. In Table 1, we indicated the best filtering runs and their input parameters. As criterion for selecting the best runs we used the $F$-value, our smoothing quality control parameter defined in Sections 4 and 5, requiring $F$ to be large. Thereafter, we confirmed the selection by visually inspecting the filtered output data cube. Table 1 highlights, for our sample, that finding the

input parameters of the best runs is not straightforward. In fact, the box and Gaussian kernels are highly dependent on the spatial and velocity extents, and the signal-to-noise ratio of the unknown faint signal. Note that the Gaussian smoothing gives better results than the box smoothing, because a gentler smoothing preserves better the shape of the data (the differences are clearly visible in the second and third panels in Fig. 4). Two examples which suffer from these limitations are:

1) ModelB: very faint signal (signal-to-noise ratio $\sim 1$) with limited extent;
2) NGC 4111: very extended, relatively faint, signal.

In the first case large kernels are necessary to considerably enhance the very low level signal. Large kernels (e.g. for the box filter $N_i > 5$ and for the Gaussian filter $FWHM_i > 3$) will, however, wash out the signal because it is not coherent at such large scales. In the second case, very large kernels ($N_i = 9$ for the box filters and $FWHM_i = 7$ for the Gaussian filter) provide the best smoothing and the maximum $F$-values. Such kernels drastically reduce, however, the spatial and velocity resolution of the data.

The optimal dimensions of the box and Gaussian kernels strongly depend on the extent of the signal and the signal-to-noise ratio. The quite different, best input parameters of ModelA, ModelB and ModelC, with their different signal-to-noise ratios, illustrate this clearly. For example, the best runs for modelB use larger kernels in the $y$ direction compared to the other models. The optimal kernels for smoothing ModelA and ModelC have, on the other hand, a very narrow $z$ component. This is expected as a higher noise level hides the signal and modifies the overall shape of the signal itself (i.e. the faintest parts will disappear into the noise).

Second, we analyzed wavelet filters in detail. Our investigation focused on thresholding the data in the wavelet domain. We performed the filtering operation exploiting a wavelet lifting algorithm. Two main wavelets have been used: the Haar and the Le Gall wavelet. Wavelet lifting is a powerful technique, but unfortunately it generates artifacts undesirable for our visualization purposes (see Fig. 5). The filtering results give very high values of the $F$-parameter as shown in Table 1. The wavelet thresholding filter, however, requires a thorough investigation of the main parameters (choice of the basic wavelet, maximum number of levels for wavelet decomposition, thresholding values for each decomposition level) for obtaining an optimal denoising of the data. We consider this a drawback for user-friendly visualization purposes.

The optimal input parameters reported in Table 1 vary for each data cube of our sample. The thresholds parameters, $t_{l,wavelet}$, have strong dependencies on the choice of the wavelet and the signal-to-noise ratio of the faint signal. Moreover, the choice of the optimal wavelet and decomposition level, $l$, depends on the extent of the faint structure. For example, the arc-shape structure in the ModelB is very thin along the velocity direction (few channels). Therefore, a the Haar wavelet and $l = 2$ are the optimal choice, while the Le Gall wavelet and a higher decomposition level, $l = 3$, provide the optimal filtering results for WEIN069, NGC3379 and NGC4111, because these data shows a more extended component.

Filtering with a higher order wavelet than Le Gall may give optimal results without requiring a pre-smoothing step. However, we showed that the choice of the wavelet is constrained by the unknown extent of the faint signal. For example, very high-order wavelets are not optimal for filtering the models.

Using different decomposition levels in each spatial and velocity dimension (or a tree structure, e.g. `Octree`; Laboratory and Meagher, 1980) may also improve the filtering quality. However, in the case of morphological complex resolved galaxies this approach is rather difficult. For example, it is necessary to determine the optimal levels of decomposition for each dimension and these depend on the signal extent and signal-to-noise ratio as well. This is analogous to the issue of finding the optimal kernel for the box and Gaussian filters.

Applying wavelet decomposition and thresholding the approximation bands, as shown in Section 3.4, is effectively a segmentation of the data. Though efficient, the disadvantage is that it also eliminates very low signal-to-noise emission if the thresholding parameters are not properly tuned to the data. Since our aim is to couple filtering techniques to visualization, thresholding techniques are not favored as they limit the interactive visual data exploration.

Third, we implemented a modification of the diffusion filter: the intensity-driven gradient filter (see 3.3). This smoothing algorithm has adaptive characteristics which helps in preserving the smaller scale structure of the signal, thus avoiding the limitations of the box and Gaussian filters. The parameters of intensity-driven gradient filter mainly depend on the signal-to-noise ratio of the emission, which we found to be quite similar for the objects studied here. In fact, the intensities of the majority of the voxels of the faint signal are between 1 and 2 $rms$. For example, in Section 2, we illustrated 3-D visualizations of the output of the intensity-driven gradient filter with default parameters ($K = 1.5$, $\tau = 0.0325$, $n = 20$ and $C_{x,y,z} = 5$) for two very different objects (WEIN069 and NGC3379). In both cases, the smoothing is successful in bringing out the low signal-to-noise structures. In fact, in the case of the gradient filter, the $F$-values of the best runs, reported in Table 1, do not differ more than 15% from the runs with default parameters.

The main input parameters ($K$, $\tau$ and $n$) of the best filtering results for the three models in Table 1 do not vary. The peak signal-to-noise ratio of ModelC is $\sim 3$ times higher than that of ModelA. Therefore, the dependencies of the input parameters of the intensity-driven gradient respect to the signal-to-noise ratio are not stiff functions.

We conclude that the intensity-driven gradient is the most promising filter because it preserves the detailed struc-

ture of the signal with high signal-to-noise ratio ($> 3$) at the highest resolution, while smoothing only the faint part of the signal (S/N $< 3$). Moreover, the input parameters need only minimal tuning to the signal itself.
On the other hand, this filter applies a diffusion process which has the following drawbacks:

a) the flux scale is not conserved and depends on the signal-to-noise ratio and hence degree of 'smoothing' or resulting resolution;
b) setting too high values of the parameters $n$ and $\tau$ can create unrealistic web structures (negative and positive) between the peaks of the negative and positive parts of the noise.

The first issue is not a problem for visualization. In fact, the main purpose of the filtering operation, in this context, is to find and enhance low-level signals. Quantitative analysis, such as calculating column densities, intensity weighted mean velocities, velocity dispersions etc., can always be performed on the original data cube once the volume that contains all the signal has been identified. Regarding the second issue: in Fig. 9 we show as a guideline the dependencies of the $F$-parameter on the input parameters $K$, $\tau$ and $n$.

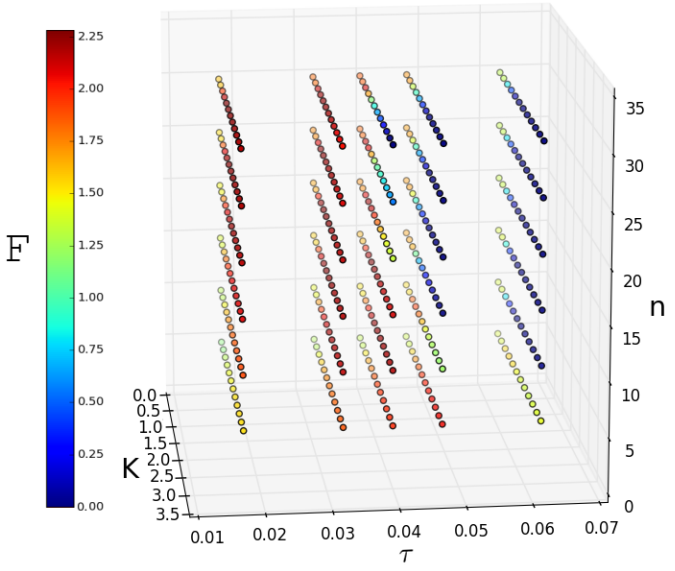


Figure 9: The $F$-values applying to WEIN069 an intensity-driven gradient filter with parameters $K$, $\tau$, $n$ and $C_{x,y,z} = 5$. In this 3-D scatter plot, the $F$-values are displayed as a 4-th dimension using a color scale. The red dots represents filtering with an high value of the parameter $F$ ($F$-values $> 1.75$). The $F$-parameter shows low values ($< 1$) for high values of $n$ and $\tau$ ($n > 15$ and $\tau > 0.0475$). For more information regarding the $F$-parameter refer to sections 4 and 5.

Finally, the previous results suggest that intensity-driven gradient smoothing can be employed for finding H I sources as well. This technique could be an alternative for the smooth-and-clip method and has the advantage that the user does not have to specify the smoothing kernels. The robustness of such a method should be tested on a larger number of different cases than we have used here. This is beyond the scope of the present investigation.

In Section 6, we reported the benchmark of our CPU and GPU implementations of the filtering algorithms investigated in this paper. The codes are publicly available[2] and we integrated them in a module of `SlicerAstro`[3], a first design of an astronomical extension of `3DSlicer`[4] (Fedorov et al., 2012). We showed that for data cubes with a number of voxels up to $5 \times 10^6$, GPU implementations of the smoothing filters can reach interactive performance (maximum execution time, $T_p < 0.3$ $s$) exploiting a GTX860M, i.e. a GPU suitable for gaming, found on laptops with mid-level performance. For data cubes up to $10^8$ voxels, the filters can still reach relatively fast performance (maximum execution time with a GTX860M, $T_p < 3.5$ $s$).

In conclusion, the GPU implementation of the intensity-driven gradient filter satisfies our filtering and visualization requirements best. The filter provides interactive performance, requires minimal tuning of the input parameters, and efficiently enhances faint structures in our data sample without degrading the resolution of the high signal-to-noise data.

## 8. Acknowledgments

We thank M.A. Ramatsoku, E. Busekool and M.A.W. Verheijen for proving us the H I data of WEIN069 and NGC4111. We thank S. Pieper (Isomics) and K. Martin (Kitware) of the 3D-Slicer project for their support with the GPU implementation of the filters. We thank M. Papastergis and A. Marasco for their feedback regarding Section 5. Finally, we thank the reviewers for their constructive comments, which helped us to improve the paper substantially.

3-D visualization screenshots shown in this paper were generated by using `3DSlicer`.

D. Punzo and J.M van der Hulst acknowledge support from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement nr. 291-531.

## 9. References

## References

J. M. van der Hulst, T. S. van Albada, R. Sancisi, The Westerbork HI Survey of Irregular and Spiral Galaxies, WHISP, in: J. E. Hibbard, M. Rupen, J. H. van Gorkom (Eds.), Gas and Galaxy Evolution, Vol. 240 of Astronomical Society of the Pacific Conference Series, 2001, p. 451.

[2] `https://github.com/Punzo/SlicerAstro/AstroSmoothing`
[3] `http://wiki.slicer.org/slicerWiki/index.php/Documentation/Nightly/Extensions/SlicerAstro`
[4] `3DSlicer` (`https://www.slicer.org/`) is a medical visualization package with advanced 3-D visualization capabilities.

A. Popping, R. Jurek, T. Westmeier, P. , L. Flöer, M. Meyer, B. Koribalski, Comparison of Potential ASKAP HI Survey Source Finders, Publications of the Astronomical Society of Australia 29 (2012) 318–339.

S. Johnston, R. Taylor, M. Bailes, N. Bartel, C. Baugh, M. Bietenholz, C. Blake, R. Braun, J. Brown, S. Chatterjee, J. Darling, A. Deller, R. Dodson, P. Edwards, R. Ekers, S. Ellingsen, I. Feain, B. Gaensler, M. Haverkorn, G. Hobbs, A. Hopkins, C. Jackson, C. James, G. Joncas, V. Kaspi, V. Kilborn, B. Koribalski, R. Kothes, T. Landecker, A. Lenc, J. Lovell, J.-P. Macquart, R. Manchester, D. Matthews, N. McClure-Griffiths, R. Norris, U.-L. Pen, C. Phillips, C. Power, R. Protheroe, E. Sadler, B. Schmidt, I. Stairs, L. Staveley-Smith, J. Stil, S. Tingay, A. Tzioumis, M. Walker, J. Wall, M. Wolleben, Science with ASKAP. The Australian square-kilometre-array pathfinder, Experimental Astronomy 22 (2008) 151–273.

A. R. Duffy, M. J. Meyer, L. Staveley-Smith, M. Bernyk, D. J. Croton, B. S. Koribalski, D. Gerstmann, S. Westerlund, Predictions for ASKAP neutral hydrogen surveys, Monthly Notices of the Royal Astronomical Society 426 (2012) 3385–3402.

M. Verheijen, T. Oosterloo, G. Heald, W. van Cappellen, HI Surveys with APERTIF, in: Panoramic Radio Astronomy: Wide-field 1-2 GHz Research on Galaxy Evolution, 2009, p. 10.

M. T. Whiting, DUCHAMP: a 3D source finder for spectral-line data, Monthly Notices of the Royal Astronomical Society 421 (2012) 3242–3256.

P. Serra, T. Westmeier, N. Giese, R. Jurek, L. Flöer, A. Popping, B. Winkel, T. van der Hulst, M. Meyer, B. S. Koribalski, L. Staveley-Smith, H. Courtois, SOFIA: a flexible source finder for 3D spectral line data, Monthly Notices of the Royal Astronomical Society 448 (2015) 1922–1929.

R. Sancisi, F. Fraternali, T. Oosterloo, T. van der Hulst, Cold gas accretion in galaxies, The Astronomy and Astrophysics Review 15 (2008) 189–223.

D. Punzo, J. M. van der Hulst, J. B. T. M. Roerdink, T. A. Oosterloo, M. Ramatsoku, M. A. W. Verheijen, The role of 3-D interactive visualization in blind surveys of H I in galaxies, Astronomy and Computing 12 (2015) 86–99.

J. M. van der Hulst, J. P. Terlouw, K. G. Begeman, W. Zwitser, P. R. Roelfsema, The Groningen Image Processing SYstem, GIPSY, in: D. M. Worrall, C. Biemesderfer, J. Barnes (Eds.), Astronomical Data Analysis Software and Systems I, Vol. 25 of Astronomical Society of the Pacific Conference Series, 1992, p. 131.
URL `https://www.astro.rug.nl/~gipsy/articles/gipsypaper.html`

M. G. R. Vogelaar, J. P. Terlouw, The Evolution of GIPSY—or the Survival of an Image Processing System, in: F. R. Harnden, Jr., F. A. Primini, H. E. Payne (Eds.), Astronomical Data Analysis Software and Systems X, Vol. 238 of Astronomical Society of the Pacific Conference Series, 2001, p. 358.

M. A. W. Verheijen, Galaxy evolution in dense environments: a concise HI perspective, in: A. Diaferio (Ed.), IAU Colloq. 195: Outskirts of Galaxy Clusters: Intense Life in the Suburbs, 2004, pp. 394–400.

E. Busekool, M. A. W. Verheijen, et al., in prep.

P. Serra, T. Oosterloo, R. Morganti, K. Alatalo, L. Blitz, M. Bois, F. Bournaud, M. Bureau, M. Cappellari, A. F. Crocker, R. L. Davies, T. A. Davis, P. T. de Zeeuw, P.-A. Duc, E. Emsellem, S. Khochfar, D. Krajnović, H. Kuntschner, P.-Y. Lablanche, R. M. McDermid, T. Naab, M. Sarzi, N. Scott, S. C. Trager, A.-M. Weijmans, L. M. Young, The ATLAS$^{3D}$ project - XIII. Mass and morphology of H I in early-type galaxies as a function of environment, Monthly Notices of the Royal Astronomical Society 422 (2012) 1835–1862.

M. Ramatsoku, M. A. W. Verheijen, R. C. Kraan-Korteweg, G. I. G. Józsa, A. C. Schröder, T. H. Jarrett, E. C. Elson, W. van Driel, W. J. G. de Blok, P. A. Henning, The WSRT ZoA Perseus-Pisces filament wide-field H I imaging survey - I. H I catalogue and atlas, Monthly Notices of the Royal Astronomical Society 460 (2016) 923–941.

R. Weinberger, W. Saurer, R. Seeberger, Penetrating the "zone of avoidance". I. A compilation of optically identified extragalactic objects within $|b| <= 5deg.$, Astronomy and Astrophysics Supplement 110 (1995) 269.

A. Goyal, A. Bijalwan, M. K. Chowdhury, A comprehensive review of image smoothing techniques, International Journal of Advanced Research in Computer Engineering & Technology 1 (4).

A. Buades, B. Coll, J.-M. Morel, A review of image denoising algorithms, with a new one, Multiscale Modeling & Simulation 4 (2) (2005) 490–530.

R. C. Gonzalez, R. E. Woods, Digital image processing.

A. R. Weeks, Fundamentals of electronic image processing, SPIE Optical Engineering Press Bellingham, 1996.

P. Serra, R. Jurek, L. Flöer, Using Negative Detections to Estimate Source-Finder Reliability, Publications of the Astronomical Society of Australia 29 (2012) 296–300.

J.-L. Starck, F. Murtagh, Image restoration with noise suppression using the wavelet transform, Astronomy and Astrophysics 288 (1994) 342–348.

A. M. Hopkins, C. J. Miller, A. J. Connolly, C. Genovese, R. C. Nichol, L. Wasserman, A New Source Detection Algorithm Using the False-Discovery Rate, The Astronomical Journal 123 (2002) 1086–1094.

E. Carlinet, T. Géraud, A fair comparison of many max-tree computation algorithms (extended version of the paper submitted to ismm 2013, arXiv preprint arXiv:1212.1819.

P. Teeninga, U. Moschini, S. C. Trager, M. H. Wilkinson, Improved detection of faint extended astronomical objects through statistical attribute filtering (2015a) 157–168.

M. A. Westenberg, J. B. T. M. Roerdink, M. H. F. Wilkinson, Volumetric attribute filtering and interactive visualization using the max-tree representation, Image Processing, IEEE Transactions on 16 (12) (2007) 2943–2952.

P. Teeninga, U. Moschini, S. C. Trager, M. H. Wilkinson, Improving background estimation for faint astronomical object detection, in: Image Processing (ICIP), 2015 IEEE International Conference on, 2015b, pp. 1046–1050.

U. Moschini, P. Teeninga, M. H. F. Wilkinson, N. Giese, D. Punzo, J. M. van der Hulst, S. C. Trager, Towards better segmentation of large floating point 3d astronomical data sets: first results, in: Proceedings of the 2014 Conference on Big Data from Space (BiDS 2014), 2014, pp. 232–235.

C. Arnoldus, A max-tree-based astronomical source finder, Master's thesis, Rijksuniversiteit Groningen, the Netherlands (2015).

P. Perona, J. Malik, Scale-space and edge detection using anisotropic diffusion, Pattern Analysis and Machine Intelligence, IEEE Transactions on 12 (7) (1990) 629–639.

S. Mallat, A wavelet tour of signal processing, Academic press, 1999.

I. Daubechies, W. Sweldens, Factoring wavelet transforms into lifting steps, Journal of Fourier analysis and applications 4 (3) (1998) 247–269.

D. L. Gall, A. Tabatabai, Sub-band coding of digital images using symmetric short kernel filters and arithmetic coding techniques, in: Acoustics, Speech, and Signal Processing, 1988. ICASSP-88., 1988 International Conference on, IEEE, 1988, pp. 761–764.

A. M. Wink, J. B. T. M. Roerdink, Denoising functional mr images: a comparison of wavelet denoising and gaussian smoothing, Medical Imaging, IEEE Transactions on 23 (3) (2004) 374–387.

L. Flöer, B. Winkel, Publications of the Astronomical Society of Australia, Astronomical Society of the Pacific Conference Series 29 (2012) 244–250.

S. Bergner, M. Sedlmair, T. Möller, S. N. Abdolyousefi, A. Saad, Paraglide: Interactive parameter space partitioning for computer simulations, IEEE Trans. Vis. Comput. Graph. 19 (9) (2013) 1499–1512.

T. S. Yoo, M. J. Ackerman, W. E. Lorensen, W. Schroeder, V. Chalana, S. Aylward, D. Metaxas, R. Whitaker, Engineering and algorithm design for an image processing api: a technical report on itk-the insight toolkit, Studies in health technology and informatics (2002) 586–592.

F. Walter, E. Brinks, W. J. G. de Blok, F. Bigiel, R. C. Kennicutt, Jr., M. D. Thornley, A. Leroy, THINGS: The H I Nearby Galaxy

Survey, The Astronomical Journal 136 (2008) 2563–2647.

W. J. v. d. Laan, A. C. Jalba, J. B. T. M. Roerdink, Accelerating wavelet lifting on graphics hardware using cuda, Ieee transactions on parallel and distributed systems 22 (1) (2011) 132–146, relation: http://www.rug.nl/informatica/onderzoek/bernoulli Rights: University of Groningen, Johann Bernoulli Institute for Mathematics and Computer Science.

T.-T. Wong, C.-S. Leung, P.-A. Heng, J. Wang, Discrete wavelet transform on consumer-level graphics hardware., IEEE Trans. Multimedia 9 (3) (2007) 668–673.

C. Tenllado, J. Setoain, M. Prieto, L. Pinuel, F. Tirado, Parallel implementation of the 2d discrete wavelet transform on graphics processing units: Filter bank versus lifting., IEEE Trans. Parallel Distrib. Syst. 19 (3) (2008) 299–310.

R. P. I. I. P. Laboratory, D. Meagher, Octree Encoding: a New Technique for the Representation, Manipulation and Display of Arbitrary 3-D Objects by Computer, 1980.
URL `https://books.google.nl/books?id=CgRPOAAACAAJ`

A. Fedorov, R. Beichel, J. Kalpathy-Cramer, J. Finet, J.-C. Fillion-Robin, S. Pujol, C. Bauer, D. Jennings, F. Fennessy, M. Sonka, J. Buatti, S. Aylward, J. Miller, S. Pieper, R. Kikinis, 3d slicer as an image computing platform for the quantitative imaging network 30 (9) (2012) 1323–41.